

CLAIMS

1 1. A method in a computer system for preparing a task to be
2 swapped out from processor utilization, the computer system having multiple
3 processors and an operating system, each processor having multiple streams for
4 simultaneously executing threads of the task, the task having one or more teams of
5 threads, each team representing threads executing on a single processor, the method
6 comprising:

7 raising an exception for each stream of each processor currently
8 executing a thread of the task; and

9 in response to the raising of the exception, for each stream executing a
10 thread,

11 saving the state of the stream;

12 determining whether the stream is a team master stream;

```
13     when the stream is not the team master stream, quitting the
14     stream;
```

15 when the stream is the team master stream,

16 waiting for all other streams executing threads in the
17 same team to quit;

18 determining whether the stream is a task master stream,
19 when the stream is not the task master stream, notifying
20 the operating system that the team for this processor is ready to be swapped out;

21 when the stream is the task master stream,

22 waiting for all other teams to notify the operating
23 system that the team is ready to be swapped out; and

24 notifying the operating system that the task is
25 ready to be swapped out.

1 2. A method in a computer system for preparing a task to be
2 swapped out from processor utilization, the computer system having multiple
3 processors and an operating system, each processor having multiple streams for
4 executing threads of the task, the task having one or more teams of threads, each
5 team representing threads executing on a single processor, the method comprising:

6 for each team, designating one stream that is executing a thread as a
7 team master stream;

8 designating one stream that is executing a thread as a task master
9 stream for the task;

10 for each team master stream, notifying the operating system that the
11 team is ready to be swapped out when each other thread of the team has quit its
12 stream; and

13 for the task master stream, notifying the operating system that the task
14 is ready to be swapped when each of the other teams have notified the operating
15 system that that team is ready to be swapped out.

1 3. The method of claim 2 wherein the operating system swaps out
2 the task upon receiving the notification that the task is ready to be swapped out.

1 4. The method of claim 2 wherein each stream stores its own state
2 before quitting the stream or notifying the operating system.

1 5. The method of claim 2 wherein each stream that is not a team
2 master stream quits its stream.

1 6. The method of claim 2 wherein the notifying of the operating
2 system by the task master stream includes indicating whether the task is blocked so

3 that the operating system can defer swapping in the task until an event occurs to
4 unblock the task.

1 7. The method of claim 2 wherein each team master stream
2 notifies the operating system of the number of streams that were executing threads
3 so that the operating system can defer swapping in the task until enough streams are
4 available to execute each of the threads that were executing when the task was
5 swapped out.

1 8. A method in a computer system for preparing a task to be
2 swapped out from processor utilization, the computer system having a processor and
3 an operating system, the processor having multiple streams for executing threads of
4 the task, the method comprising:

5 designating one stream that is executing a thread as a master stream;
6 saving the state of each stream that is executing a thread;
7 under control of each stream that is not the master stream, quitting the
8 stream;

9 under control of the master stream, notifying the operating system that
10 the task is ready to be swapped out.

1 9. The method of claim 8 wherein the operating system swaps out
2 the task upon receiving the notification that the task is ready to be swapped out.

1 10. The method of claim 8 wherein each stream saves its own state.

1 11. The method of claim 8 including upon being swapped in:
2 starting execution of a master stream; and
3 under control of the master stream,

4 creating a stream corresponding to each stream that quit when
5 the task was swapped out; and

6 initializing each created stream based on state information
7 saved before the stream quit.

1 12. The method of claim 8 wherein the notifying of the operating
2 system by the master stream includes indicating whether the task is blocked so that
3 the operating system can defer swapping in the task until an event occurs to unblock
4 the task.

1 13. The method of claim 8 wherein the task notifies the operating
2 system of the number of streams that were executing threads so that the operating
3 system can defer swapping in the task until enough streams are available to execute
4 each of the threads that were executing when the task was swapped out.

1 14. A method in a computer system for scheduling tasks, the
2 method comprising:

1 15. The method of claim 14 wherein the computer system is a
2 multithreaded computer system.

1 16. The method of claim 14 wherein in response to the notification,
2 the task saves its state.

1 17. A method in a computer system for restarting the execution of a
2 task that has been swapped out of processor utilization, the method comprising:
3 starting execution of a master stream; and
4 under control of the master stream,
5 creating a stream corresponding to each stream that quit when
6 the task was swapped out; and
7 initializing each created stream based on state information
8 saved before the stream quit.

1 18. A method in a computer system for returning a stream to a task
2 executing an operating system call that is blocked, the computer system having a
3 processor with multiple streams, method comprising:
4 under control of the operating system executing on a stream, invoking
5 a function provided by the task;
6 under control of the invoked function, executing instructions of the
7 task on that stream; and
8 under control of the operating system, notifying the task when the
9 operating system call is complete.

1 19. The method of claim 18 wherein the notifying includes
2 invoking a function provided by the task using a stream of the
3 operating system; and
4 under control of that invoked function,
5 indicating that the operating system call is complete; and

6 invoking another operating system call to return the operating
7 system stream to the operating system.

3 indicating that a thread that invoked the operating system call is
4 blocked; and

5 executing another thread on that stream.

1 21. A method in a computer system for performing an inter-thread
2 long jump, the method comprising:

3 receiving an indication of a set jump location;

4 determining whether the set jump thread that set the set jump location
5 is the same thread that is currently executing; and

when the set jump thread is not the same thread that is currently executing, setting the state of the set jump thread to execute a long jump indicating the set jump location.

1 22. The method of claim 21 wherein the setting of the state
2 includes setting the program counter of the set jump thread to point to a long jump
3 routine.

1 23. The method of claim 21 wherein when the set jump thread is
2 blocked on an operating system call aborting the operating system call.

1 25. The method of claim 21 wherein when the set jump thread is
2 running notifying the set jump thread of the inter-thread long jump.

1 26. The method of claim 21 wherein when the set jump thread is
2 running directing the set jump thread to enter a known state prior to setting the state
3 of the set jump thread.

1 27. The method of claim 26 wherein the known state is a quiescent
2 state.

1 28. The method of claim 21 including determining whether the set
2 jump thread still exists.

1 29. The method of claim 21, 22, 23, 24, 25, 26, 27, or 28 wherein
2 the computer system supports multiple streams.

1 30. A method in a computer system for a server to coordinate
2 assignment of a resource to clients, the method comprising:

3 assigning the resource to a client;

4 receiving notification from the client assigned to the resource that the
5 client is waiting for an occurrence of an event before the resource can be
6 productively used;

7 upon receiving the notification, unassigning the resource from the
8 client; and

9 after occurrence of the event, reassigning the resource to the client.

1 31. The method of claim 30 wherein the server upon receiving the
2 notification assigns the resource to another client.

1 32. The method of claim 30 wherein the server is an operating
2 system, the clients are tasks, and the resource is a processor resource of the
3 computer system.

1 33. The method of claim 32 wherein the processor resource is
2 single-threaded.

1 34. The method of claim 32 wherein the processor resource is
2 multi-threaded.

1 35. The method of claim 32 wherein the operating system upon
2 notification assigns the processor resource to another task.

1 36. The method of claim 32 wherein the task can only perform idle
2 processing until the event occurs.

1 37. The method of claim 32 wherein the task can perform no
2 processing until the event occurs.

1 38. The method of claim 32 wherein the operating system reassigns
2 the processor resource to the task whenever an external event directed to the task
3 occurs.

1 39. The method of claim 32 wherein the notification is received by
2 the operating system in response to informing the task that the task is to be swapped
3 out from processor resource utilization.

1 40. The method of claim 32 wherein the operating system assigns a
2 task to the processor resource by assigning the task to a domain.

1 41. A method in a computer system for clients to coordinate
2 assignment of a resource with a server, the method comprising:
3 determining whether the client cannot productively use the resource
4 until an event occurs; and

5 when the client determines that it cannot productively use the
6 resource, notifying the server that the client cannot productively use the resource
7 wherein the server unassigns the resource from the client until after an
8 event occurs.

1 42. The method of claim 41 wherein the server is an operating
2 system, the clients are tasks, and the resource is a processor resource of the
3 computer system.

1 43. The method of claim 42 wherein the processor resource is
2 single-threaded.

1 44. The method of claim 42 wherein the processor resource is
2 multi-threaded.

1 45. The method of claim 42 wherein the operating system upon
2 notification assigns the processor resource to another task.

1 46. The method of claim 42 wherein the task can only perform idle
2 processing until the event occurs.

1 47. The method of claim 42 wherein the task determines whether it
2 cannot productively use the processor resource until an event occurs after receiving
3 an indication from the operating system that the task is to be swapped out from
4 processor resource utilization.

1 48. The method of claim 42 wherein the task notifies the server
2 after receiving an indication from the operating system that the task is to be swapped
3 out from processor resource utilization.

1 49. The method of claim 42 wherein the task can perform no
2 processing until the event occurs.

1 50. The method of claim 42 wherein the operating system assigns a
2 task to the processor resource by assigning the task to a domain.

1 51. A computer-readable medium for causing an operating system
2 executing on a computer to coordinate assignment of processor resources to task, by:
3 assigning a processor resource to a task, the task having productive
4 instructions to execute using the assigned processor resource;

5 receiving notification from the task assigned to the processor resource
6 that the task is waiting for an occurrence of an event before the task can continue
7 execution of the productive instructions

8 upon receiving the notification, unassigning the processor recourse
9 from the task; and

10 after occurrence of the event, reassigning the processor resource to the
11 task.

1 52. The computer-readable medium of claim 51 wherein the
2 operating system upon receiving the notification assigns the processor resource to
3 another task.

1 53. The computer-readable medium of claim 51 wherein the
2 processor is single-threaded.

1 54. The computer-readable medium of claim 51 wherein the
2 processor is multi-threaded.

1 55. The computer-readable medium of claim 51 wherein the task
2 can only perform idle processing until the event occurs.

1 56. The computer-readable medium of claim 51 wherein the task
2 can perform no processing until the event occurs.

1 57. The computer-readable medium of claim 51 wherein the
2 operating system reassigns the processor resource to the task after an external event
3 directed to the task occurs.

1 58. The computer-readable medium of claim 51 wherein the
2 notification is received by the operating system in response to informing the task
3 that the task is to be swapped out from processor resource utilization.

1 59. A computer-readable medium for causing clients executing on a
2 computer system to coordinate assignment of a processor resource with an operating
3 system, by:

4 under control of a task,

5 determining whether the task cannot productively use the
6 processor resource until an event occurs; and

7 when the task determines that it cannot productively use the
8 processor resource, notifying the operating system that the task cannot productively
9 use the resource

10 wherein the operating system unassigns the resource from the task
11 until after an event occurs.

1 60. The computer-readable medium of claim 59 wherein the
2 processor is single-threaded.

1 61. The computer-readable medium of claim 59 wherein the
2 processor is multi-threaded.

1 62. The computer-readable medium of claim 59 wherein the
2 operating system upon notification assigns the processor resource to another task.

1 63. The computer-readable medium of claim 59 wherein the task
2 can only perform idle processing until the event occurs.

1 64. The computer-readable medium of claim 59 wherein the task
2 determines whether it cannot productively use the processor resource until an event
3 occurs after receiving an indication from the operating system that the task is to be
4 swapped out from processor utilization.

1 65. The computer-readable medium of claim 59 wherein the task
2 notifies the operating system after receiving an indication from the operating system
3 that the task is to be swapped out from processor utilization.

1 66. The computer-readable medium of claim 59 wherein the task
2 can perform no processing until the event occurs.

1 67. A method in a computer system for placing a task in a known
2 state, the task having multiple threads executing on streams of the computer system,
3 the method comprising:

4 notifying each of the threads of the task executing on a stream to enter
5 a known state; and

6 for each of the threads,

7 in response to receiving the notification, entering the known so
8 that so that an action can be performed with the task being in the known state.

1 68. The method of claim 67 wherein the known state is a quiescent
2 state.

1 69. The method of claim 67 wherein the known state is where each
2 of the threads is executing idle instructions.

1 70. The method of claim 67 wherein the known state is where the
2 threads stop executing instructions.

1 71. The method of claim 67 wherein the task is assigned to a
2 protection domain and the notifying including raising a domain signal for the
3 protection domain.

1 72. The method of claim 67 wherein prior to entering the known
2 state each thread saves its state.

1 73. The method of claim 67 wherein a thread of the task initiates
2 the notifying.

1 74. The method of claim 73 wherein the thread initiates the
2 notifying by sending a request to an operating system.

1 75. The method of claim 73 wherein the task notifies an operating
2 system that the task is blocked from further productive use of a resource until an
3 event occurs prior to entering the known state.

1 76. The method of claim 67 wherein the action to be performed is
2 swapping out the task from processor utilization.

1 77. The method of claim 76 wherein prior to entering the known
2 state each thread saves its state information.

1 78. The method of claim 67 wherein each thread of the task has
2 state information and the action to be performed is review of the state information.

1 79. The method of claim 78 wherein the review of the state
2 information is by a debugger.

1 80. The method of claim 79 wherein the debugger executes as a
2 thread of the task that does not enter the known state.

1 81. The method of claim 67 wherein the action is to process signal
2 indicated by an operating system.

1 82. The method of claim 67 wherein the action is to perform an
2 inter-thread long jump.

1 83. The method of claim 67 wherein the known state is waiting on
2 a synchronization indication.

1 84. The method of claim 83 wherein the waiting is performed by
2 accessing a memory location with a synchronization access mode of future.

1 85. The method of claim 67 wherein entering the known state
2 includes invoking an operating system call.

1 86. The method of claim 67 wherein a master thread is designated
2 to notifying an operating system when the task is in the known state.

1 87. The method of claim 67 wherein one of the threads enters a
2 known state of processing signals while the other threads are in a known state that is
3 quiescent state.

1 88. The method of claim 87 wherein after processing the signals
2 each of the threads exit the known state.

1 89. A method in a computer system for a task to exit a known state,
2 the computer system supporting multiple streams, the task having multiple threads
3 executing on the streams, the method comprising:

4 notifying each of the threads of the task to exit the known state; and
5 for each thread,

6 in response to receiving the notification, executing instructions
7 that were to be executed prior to entering the known state.

1 90. The method of claim 89 wherein the known state is a quiescent
2 state.

1 91. The method of claim 89 wherein the known state is where the
2 threads are executing idle instructions.

1 92. The method of claim 89 wherein the known state is where the
2 threads stop executing instructions.

1 93. The method of claim 89 wherein the known state is where the
2 threads are waiting on a synchronization indication.

1 94. The method of claim 93 wherein the notifying includes
2 indicating the synchronization.

1 95. The method of claim 89 wherein each thread restores state
2 information that was save prior to the state entering the known state.

1 96. The method of claim 95 wherein one thread performs signal
2 processing upon exiting the known state.

1 97. The method of claim 96 wherein the other threads wait until the
2 signals are processed before executing instructions that were to be executed prior to
3 entering the known state.

1 99. The method of claim 89 including reserving a number of
2 streams for the task.

1 100. The method of claim 89 wherein after receiving the
2 notification, the task creates streams for the threads.

1 101. A method in a computer system for assigning a processor
2 resource to a thread of a task, the method comprising:

3 under control of a thread of the task, invoking an operating system call
4 that will block waiting for the occurrence of an event;

5 under control of the operating system, invoking a routine of the task so
6 that the routine can assign the processor resource to another thread of the task.

1 102. The method of claim 101 wherein the processor resource is a
2 stream of a processor that supports multiple streams.

1 103. The method of claim 101 wherein the task registers the routine
2 with the operating prior to invoking the operating system call.

1 104. The method of claim 101 including notifying the task when a
2 operating system call completes.

1 105. The method of claim 101 wherein the computer system is a
2 MTA computer system.